


FNAL Grid Facility

July 20, 2004



Grid Facility at FNAL



There have been proposals floated for establishing a shared Grid Enabled Facility at FNAL

- ➡ There are existing Grid enabled clusters associated with specific projects (CMS, SAM-GRID)
- Designed to meet the needs to the builder and frequently running in a “prototype” mode
 - Security exemptions, high operational load, less than production quality service

A more formal grid resource at FNAL would be an interesting development effort

- ➡ What kind of Facility to build (Shared, opportunistic, schedule-able?)
- ➡ How to manage and operate production quality grid services
 - Operational model, support, security models



Starting the Open Science Grid



FNAL has the opportunity to build the first facility infrastructure for the Open Science Grid

- ➡ There are many grid services to develop on the way to a fully functional persistent grid infrastructure
 - FNAL can contribute to many of these
- ➡ The hardware and the distribution are what defines the scale of a grid
- ➡ FNAL has the opportunity and the means to deploy a grid enabled cluster which is large enough to be taken seriously
 - As a National Lab we have a natural leadership role

A large grid facility should only be built if it can benefit the needs of potential Open Science Grid stakeholders

- ➡ CMS and ATLAS have both demonstrated the ability to capitalize on distributed opportunistic computing resources
- ➡ It would be interesting to see if the Run2 program would also benefit
 - Design this in from the beginning

There is a hurdle to use a facility that you don't completely control

- ➡ Need to adjust your environment and your way of thinking

In order to make people want to expend the effort you need to make an attractive enough target

- ➡ US-CMS is proposing building the facility out of new equipment
 - The majority of the FY05 hardware procurement could be contributed to a Grid enabled farm labeled OSG
 - CMS contribution would be 250-300 dual nodes
 - It would be good to get to 1k CPUs, so another 200 duals from somewhere
 - Make a flexible enough architecture to support the many potential use patterns
 - Build in HDCF

Fermilab's Storage infrastructure is already an attractive target

- ➡ Continue to develop, improve, advertise, and use the grid interfaces to the mass storage system

Processing, storage, and network is a lot of the way there

The Open Science Grid program of work has started with a proposed blue print

- ➡ Calls for a milestones for a distributed system with a defined scale and functionality by Feb. 2005
- ➡ OSG-0 will probably look like an evolution of Grid2003
 - Standard Grid services (GRAM, Information providers, monitors, GridFTP and SRM transfers, etc)
 - We know this won't scale arbitrarily.

As this is a FNAL facility we should look at what would be needed for efficient Run2 use

- ➡ Interoperate low level grid services with SAM-Grid services on the same physical set of worker nodes
- ➡ Operate CAF type services through Condor Glide-in

The interfaces should be grid based so development proceeds to using generic distributed resources, but the facility is close and well connected to FNAL storage

Support Open Science Grid stakeholders

- ➔ US-LHC, Run2, RHIC, LIGO, SDSS, Biology, NSF Education and Outreach,....

How to partition resources?


- ➔ If CMS is the only group that provides resources, it will be an opportunistic facility.
 - CMS resources are heavily used, but the timescale of use is different from Run2, so there might be a reasonable synergy
- ➔ If other groups step up with resources, then it needs to be a shared facility where contributors get at least what they put it
 - Opportunistic for other use

What does it mean to support groups at a grid enabled facility

- ➔ In Grid2003 this implied an operational load helping people figure why their applications don't run
 - Normally done as best effort.
 - Needs to be included in effort estimates



How to manage it?



We don't really have an operational model for managing a shared computing resource

- ➡ From the CMS side it seems like a significant risk to hand off a big physical resource like 300 nodes to an untested management structure

To me it make sense to separate the architecture and farm manage from the grid service

- ➡ US-CMS operates the physical resources
 - At the very least we have a facility that meets the CMS needs as well as the existing production facilities do
- ➡ The Grid Interfaces should be handled as a common project in the context of Open Science Grid
 - Contributions from CMS, CD, stakeholders
 - Develop the tools to enable efficient use of the facility by a lot of folks
 - Develop the policy infrastructure to meet the obligations to contributors and provide opportunistic use to others

If we have a desire to drive a persistent distributed computing infrastructure for science in the US we need to appear on the radar

- ➡ Big target farm helps, though certainly it isn't enough

A big grid enabled facility that where the Grid interfaces are controlled and developed by us would provide a good development platform for several stakeholders

- ➡ You can do environment development in a situation where the access to the data performs sufficiently
 - The network between HDCF and FCC is about what CMS expects a Tier-2 will have to FNAL by the start of the experiment.
- ➡ The grid interfaces and abstraction can be applied incrementally
 - Start small, simple, and useful and then move to complicated